

DARPA Tech, DARPA's 25th Systems and Technology Symposium
August 8, 2007
Anaheim, California
Teleprompter Script for Dr. Dan Oblinger, Program Manager,
Information Processing Technology Office

Learning/Reasoning

» **DANIEL OBLINGER:**

Decision-making is a deeply cognitive task.

As we develop automated systems to assist us with rapid decision-making,

we must confront the challenge of endowing software with deep cognitive capabilities.

These capabilities will drive a new level of interaction and generality, as exemplified in this video.

Clearly, this is impressive, but what's the revolution needed to make it real?

It's putting comprehension into the software.

Notice the system understood the goal implicit in the request: that the user is trying to identify *any* patterns of activity associated with bombing events – not just the suggested one.

Secondly, the system understood the basic structure of explicitly planned action.

Having observed the effects of this coordinated activity, the system searched for how and where the planning occurred.

It understood that organized activity could be inferred from commonalities observed across the movements of individuals.

You could call this situational awareness
but it really reflects a capability that machines currently don't have –
to deeply understand user intent and the context of the information
being discussed; in other words, “situational understanding”.

Notice, the commander does not interact with the computer as a tool to
be used *by* his team,
but rather as a
member *of* his team.

The key to enabling this shift in roles is giving computers situational understanding.

Instead of belaboring the obvious value of this,
let's focus on the question of whether it's even possible and, if so,
how do we proceed?

Well the research has already begun and

we are making progress in developing systems with cognitive capabilities.

For example,

the PAL effort is developing cognitive agents that can integrate diverse knowledge.

This is required to help warfighters organize and relate information by learning their priorities and constraints.

These technologies are being readied for demonstrations in systems that are
in the field.

We also have a number of initiatives focused on natural forms of instruction.

One effort involves the development of a system that can learn to execute complex planning tasks simply by being shown the tasks being performed.

We are also working on a general-purpose technology that can be embedded in any computing system so that system can learn from instruction.

Such technology will enable humans to instruct deployed systems by using natural human instruction, such as showing examples, explaining failures, and building upon previously learned knowledge.

These pursuits are important because they advance the deep situational understanding of our computing systems.

Situational understanding requires capturing the vast quantities of information specific to each situation.

Since that information is only known to teammates in the field, not programmers tucked safely away in some cubicle, it strongly suggests instructional computing as the key ingredient for situational understanding.

In the following vignette, we consider how this situational understanding might arise from real-time instruction.

In this example, a city planner is trying to develop effective evacuation routes, and wants to constrain the analysis using knowledge of bottlenecks that occur in day-to-day traffic.

It occurs to her, that traffic cam data might be co-opted for this task.

Since the traffic cam data base was not intended for this purpose, she must bootstrap capabilities and concepts from more primitive, hard-coded ones.

Notice that initially the system wasn't even aware of bottlenecks, much less the difference between acute and chronic ones.

It learned these concepts on the fly.

Many current IPTO programs are making great strides toward trainable systems that can acquire these kinds of deep understanding.

But even when we achieve this, we will not have replicated the plasticity of the human brain, and its ability to build upon inherently imprecise concepts.

We will ultimately need to do this if our computers are to be completely at home in our inherently imprecise world.

And there's reason to believe this is attainable.

Recent advances over multiple areas of neuroscience suggest that much of the rich and varied structure of the human neo-cortex is shaped by the inputs it receives.

Indeed, the organization of the neo-cortex *could* largely be the natural consequence of a relatively simple algorithm, adapting itself to the structure latent in the world.

Put another way,

perhaps the reason there are edge detectors in the brain is because there are edges in the world.

This suggests a tantalizing possibility:

it might be possible to construct a single, massively parallel, general-purpose algorithm that could start with

zero knowledge of its environment, but then grow to represent the structure latent in that environment by iteratively capturing each level of detail, building on the organization of the level below.

If such an algorithm were possible, it would give computing systems a radically new ability to assimilate knowledge about a situation simply by being immersed *in* that situation, much in the same way that the neo-cortex learns.

With the right

building-block algorithm,

a knowledge acquisition system could be self-architecting and capable of adapting to the inherent unpredictability of the real world.

Imbuing computing systems with such capabilities will enable them to discern regularities in their environment:

without reprogramming specifically for each context.

What should we expect of systems that marry a computer's ability to rapidly process large amounts of information with a deeper understanding of that information?

Well, imagine a world in which a commander can adapt faster than an

asymmetric adversary because his support structure is able to keep pace with his fast-changing tactics.

Today, this is infeasible, because the countless logistical changes required by a new strategy must be performed manually.

But if the computing systems managing those logistics *deeply understood* those logistics, such shifts in strategy could be instantaneous.

Or imagine a world in which our computing systems can predict not only the *physical* consequences of our actions, but the complex *social* consequences, as well.

The convergence of recent results from AI and neuroscience, along with expected gains in computational speed, has changed the playing field.

Now is the time to pursue this agenda aggressively.

With your help,
your enthusiasm,
and most importantly,
your ideas,
we *will* make it happen.

Next up is my colleague, Dr. Joe Olive,
who is going to talk to us about challenges in
Language Processing.